



# **Mellanox ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.0 Release Notes**

---

Rev 4.15.13.2

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "ASIS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies  
350 Oakmead Parkway Suite 100  
Sunnyvale, CA 94085  
U.S.A.  
[www.mellanox.com](http://www.mellanox.com)  
Tel: (408) 970-3400  
Fax: (408) 970-3403

© Copyright 2018. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, NPS®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

# Table of Contents

<b>Table of Contents</b>	<b>3</b>
<b>List Of Tables</b>	<b>4</b>
<b>Release Update History</b>	<b>5</b>
<b>Chapter 1 Overview</b>	<b>6</b>
1.1 Content of MLNX-NATIVE-ESX Driver Package	6
1.2 Supported HCAs Firmware Versions	6
1.3 Tested Hypervisors in Paravirtualized and SR-IOV Environments	7
<b>Chapter 2 Changes and New Features in Rev 4.16.13.2</b>	<b>8</b>
<b>Chapter 3 Known Issues</b>	<b>9</b>
<b>Chapter 4 Bug Fixes History</b>	<b>13</b>
<b>Chapter 5 Change Log History</b>	<b>14</b>

## List Of Tables

Table 1:	Release Update History . . . . .	5
Table 2:	Supported Uplinks to Servers . . . . .	6
Table 3:	Supported HCAs Firmware Versions . . . . .	6
Table 4:	Tested Hypervisors in Paravirtualized and SR-IOV Environments . . . . .	7
Table 5:	Changes and New Features . . . . .	8
Table 6:	Known Issues . . . . .	9
Table 7:	Fixed Bugs List . . . . .	13
Table 8:	Change Log History . . . . .	14

## Release Update History

*Table 1 - Release Update History*

Release	Date	Description
Rev 4.15.13.2	February 27, 2018	Initial release of this MLNX-NATIVE-ESXi-ConnectX-4/ConnectX-4 Lx/ConnectX-5 driver version

# 1 Overview

These are the release notes of Mellanox ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.0. Mellanox ConnectX-4/ConnectX-5 NATIVE ESXi Driver for VMware vSphere 6.0 supports the following uplinks to servers

**Table 2 - Supported Uplinks to Servers**

Version	OS	Uplink Speed
4.15.13.2	ESXi 6.0, ESXi 6.0 U1	10/40GbE
	ESXi 6.0 U2	10/25/40/50GbE



ConnectX®-4 adapter card was tested at 100Gb with ESXi 6.0 and ESXi 6.5. However, VMware has certified and will report 100GbE from ESXi 6.5 onward.

## 1.1 Content of MLNX-NATIVE-ESX Driver Package

The MLNX-NATIVE-ESX driver package is distributed as an offline bundle (.zip file) and contains:

- **ESXi 6.0:**  
MLNX-NATIVE-ESX-ConnectX-4-5\_4.15.13.2-10EM-600.0.0.2768847.zip - Hypervisor bundle for ESXi 6.0 contains the following kernel modules:
  - nmlx5\_core (Mellanox Technologies Ethernet driver)

## 1.2 Supported HCAs Firmware Versions

MLNX-NATIVE-ESX Rev 4.15.13.2 supports the following Mellanox Ethernet HCA:

**Table 3 - Supported HCAs Firmware Versions**

HCAs	Recommended Firmware Rev.
ConnectX-4	12.21.2010
ConnectX-4 Lx	14.21.2010
ConnectX-5	16.21.2010
ConnectX-5 Ex	16.21.2010

Please note that older firmware versions were not tested with this release.

For the latest firmware versions, visit:

- [http://www.mellanox.com/page/vmware\\_matrix](http://www.mellanox.com/page/vmware_matrix)
- or
- <http://www.mellanox.com/supportdownloader/>

## 1.3 Tested Hypervisors in Paravirtualized and SR-IOV Environments

**Table 4 - Tested Hypervisors in Paravirtualized and SR-IOV Environments**

Tested Hypervisors	HCAs	Operating System
SR-IOV	ConnectX-4/ConnectX-4 Lx ConnectX-5/ConnectX-5 Ex	Windows Server 2012 R2
		RedHat 6.5
		RedHat 6.6
		RedHat 7.0
		RedHat 7.1
		RedHat 7.2
		SLES 11.3
Paravirtualized	ConnectX-4/ConnectX-4 Lx ConnectX-5/ConnectX-5 Ex	Windows Server 2012 R2
		RedHat 6.5
		RedHat 6.6
		RedHat 7.0
		RedHat 7.1
		RedHat 7.2
		SLES 11.3

## 2 Changes and New Features in Rev 4.15.13.2

*Table 5 - Changes and New Features*

Feature/Change	Description
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 13</a>



### 3 Known Issues

The following is a list of general limitations and known issues of the various components of this MLNX-NATIVE-ESX release.

**Table 6 - Known Issues (Sheet 1 of 4)**

Internal Ref.	Description
1298457	<b>Description:</b> On rare occasions, a PSOD may occur when changing MTU during traffic.
	<b>Workaround:</b> Stop the traffic prior to changing the MTU configuration
	<b>Keywords:</b> MTU, PSOD
	<b>Discovered in Version:</b> 4.15.13.2
1191657	<b>Description:</b> The maximum value of RSS must be lower than the number of CPU cores.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> RSS
	<b>Discovered in Version:</b> 4.5.12.12
-	<b>Description:</b> The hardware can offload only up to 256B of headers.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Hardware offload
	<b>Discovered in Version:</b> 4.5.10.3/4.15.10.3
781277	<b>Description:</b> The "esxcli network sriovnic vf stats" command is not supported. When running this command on a vmknic, a failure message is displayed.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> esxcli SR-IOV
	<b>Discovered in Version:</b> 4.5.10.3/4.15.10.3
685558	<b>Description:</b> There is no traffic between PV and SR-IOV VF connected to different ports on the same HCA.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> PV, SR-IOV VF, HCA
858972	<b>Description:</b> Setting the "Allow Guest MTU Change" option in vSphere Client is currently not functional. Although guest MTU changes in SR-IOV are allowed, they do not affect the port's MTU and the guest's MTU remains the same as the PF MTU.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> MTU, SR-IOV
870158	<b>Description:</b> Although 'drss' and 'rss' parameters are disabled by default, when querying the nmlx5_core module parameter the displayed default values of drss/drss are "4".
	<b>Workaround:</b> N/A
	<b>Keywords:</b> DRSS, NetQ RSS

**Table 6 - Known Issues (Sheet 2 of 4)**

Internal Ref.	Description
939198	<b>Description:</b> VST mode in SR-IOV in ConnectX-5 is currently not functional.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VSNetQ RSS, encapsulated trafficT Mode, SR-IOV
1043285	<b>Description:</b> While running <code>stress ipv6 all2all traffic</code> , the MTU is changed several times and PSOD is excepted #14 (page fault).
	<b>Workaround:</b> Reboot the driver
	<b>Keywords:</b> MTU, PSOD, Traffic
1031520	<b>Description:</b> When a guest is assigned an IB PCI passthru device or an IB VF, VMware Tools networking information for the guest may be incorrect. This affects how the guest networking information, such as interfaces and their IPs, is displayed in vCenter.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VMware Tools networking information
1064844	<b>Description:</b> <code>nmlxcli</code> tool is not supported on ESXi 5.5.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> <code>nmlxcli</code> , ESXi 5.5
1064883	<b>Description:</b> Operations on <code>vmnics</code> which are in passthru mode are not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> <code>vmnics</code> , passthru mode
746100	<b>Description:</b> The <code>'esxcli mellanox uplink link info -u &lt;vmnic_name&gt;'</code> command reports the 'Auto negotiation' capability always as 'true'.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> 'Auto negotiation' capability
1068621	<b>Description:</b> SMP MADs ( <code>ibnetdiscover</code> , <code>sminfo</code> , <code>iblinkinfo</code> , <code>smpdump</code> , <code>ibqueryerr</code> , <code>ibdiagnet</code> and <code>smpquery</code> ) are not supported on the VFs.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SMP MADs
-	<b>Description:</b> Multicast and IPv6 traffic might be unstable over SR-IOV.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV, IPv6, Multicast
656416	<b>Description:</b> Reboot is required after any SR-IOV configuration change.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV

**Table 6 - Known Issues (Sheet 3 of 4)**

Internal Ref.	Description
653679	<b>Description:</b> Firmware VF configuration must be N+1 (while N is the required VF number). For example: If your configuration requires 10 VFs, the firmware must be set to support 16 VFs (ESXi Limitation).
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VFs
641252	<b>Description:</b> Driver specific statistics are not supported in ESXi 5.5.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> ESXi 5.5, Statistics
778371	<b>Description:</b> Wake-on-LAN does not notify when invalid parameters are provided.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> WoL
778572	<b>Description:</b> Nested ESXi might not function properly.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Nested ESXi
765008	<b>Description:</b> Device RSS fails to hash traffic to sufficient RX rings with Broadcast traffic.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> RSS, RX rings
852883	<b>Description:</b> In stress condition ‘Watchdog’ may appear, leading to uplink going up and down.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> uplink, watchdog
850886	<b>Description:</b> Call trace might occur after running VGT with heavy traffic.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VGT, heavy traffic
775668	<b>Description:</b> VMs can get Call Trace upon MTU change during heavy traffic.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VM, MTU, heavy traffic
776274	<b>Description:</b> Reloading the driver when the SR-IOV VFs are ON, will result in Purple Screen of Death (PSOD).
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Driver reload
784293	<b>Description:</b> VGT traffic over VXLAN interfaces is currently not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VGT traffic, VXLAN

**Table 6 - Known Issues (Sheet 4 of 4)**

Internal Ref.	Description
846359	<b>Description:</b> The adapter card might get stuck in Down state after setting the ring size to 8192.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Ring size
862168	<b>Description:</b> VMs with SR-IOV cannot be powered on when running low on available vectors.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VM, SR-IOV
777205	<b>Description:</b> Occasionally, untagged traffic can pass between VMs with SR-IOV enabled when portgroup is configured for VLAN trunk range.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> VM, SR-IOV, VLAN trunk

## 4 Bug Fixes History

Table 7 lists the bugs fixed in this release.

**Table 7 - Fixed Bugs List**

Internal Ref.	Description
1253564	<b>Description:</b> Disabled multicast loopback to avoid a scenario that prevented MAC learning in some configurations.
	<b>Keywords:</b> MAC, multicast loopback
	<b>Discovered in Release:</b> 4.15.12.12
	<b>Fixed in Release:</b> 4.15.13.2
958154	<b>Description:</b> NetQ RSS for encapsulated traffic is currently not supported. Encapsulated traffic (VXLAN/Geneve) directed to NetQ RSS queue will not be distributed through all queues' channels, thus will not utilize the RSS feature.
	<b>Note:</b> It is highly recommended to avoid requesting RSS for encapsulated interfaces, i.e. refrain from defining the following in the VM configuration file: <iface_name>.pnictFeatures=4
	<b>Keywords:</b> NetQ RSS, encapsulated traffic
	<b>Discovered in Release:</b> 4.15.8.8
	<b>Fixed in Release:</b> 4.15.12.12
698142/637104	<b>Description:</b> Traffic loss of large packets might occur after MTU change.
	<b>Keywords:</b> MTU, Traffic loss
	<b>Discovered in Release:</b> 4.15.6.22/4.5.6.22
	<b>Fixed in Release:</b> 4.15.10.3/4.5.10.3
745415	<b>Description:</b> In ESXi5.5 – when using a server with many ConnectX-4 / ConnectX-4 Lx ports, some of the interfaces do not show up in the <code>esxcfg-nics -l</code> list. This can occur if no MSI-X resources are available to enable loading all interfaces.
	<b>Keywords:</b> ESXi 5.5, port interfaces, MSI-X
	<b>Discovered in Release:</b> 4.15.4.1000/4.5.2.1000
	<b>Fixed in Release:</b> 4.15.6.22/4.5.6.22
689573	<b>Description:</b> When SR-IOV is enabled and the <code>max_vfs</code> is not equal to 0, new filters are not applied.
	<b>Keywords:</b> SR-IOV, filters
	<b>Discovered in Release:</b> 4.15.4.1000/4.5.2.1000
	<b>Fixed in Release:</b> 4.15.6.22/4.5.6.22
781255	<b>Description:</b> Added a new module parameter to <code>nmlx5_core</code> "supported_num_ports" to support up to "8" ConnectX-4 Lx ports.
	<b>Keywords:</b> Ports, module param
	<b>Discovered in Release:</b> 4.15.4.1000/4.5.2.1000
	<b>Fixed in Release:</b> 4.15.5.10/4.5.5.10

**Table 7 - Fixed Bugs List**

Internal Ref.	Description
760021	<b>Description:</b> Fixed an issue which prevented the driver from loading on machines with 64 or more CPU cores.
	<b>Keywords:</b> MSI-X
	<b>Discovered in Release:</b> 4.15.4.1000/4.5.2.1000
	<b>Fixed in Release:</b> 4.15.4.1100/4.5.4.1100

## 5 Change Log History

**Table 8 - Change Log History (Sheet 1 of 2)**

Feature/Change	Description
Rev. 4.15.12.12	
<b>SR-IOV max_vfs module parameter Type Modification</b>	Changed the type of the SR-IOV max_vfs module parameter from a single integer value to an array of unsigned integers. For further information, refer to the User Manual.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 13</a>
Rev. 4.15.10.3/4.5.10.3	
<b>ESXi CLI</b>	Added ESXi CLI support for ESXi 6.0
<b>Adapter Cards</b>	Added support for ConnectX-5/ConnectX-5 Ex adapter cards. <b>Note:</b> ConnectX-5/ConnectX-5 Ex cards are currently at beta level.
Rev. 4.15.6.22/4.5.6.22	
<b>Set Link Speed</b>	Enables you to set the link speed to a specific link speed supported by ESXi. For further information, see the User Manual section “ <i>Set Link Speed</i> ”.
<b>Priority Flow Control (PFC)</b>	Applies pause functionality to specific classes of traffic on the Ethernet link. For further information, see the User Manual section “ <i>Priority Flow Control (PFC)</i> ”.
<b>NetQ RSS</b>	Allows the user to configure multiple hardware queues backing up the single RX queue. NetQ RSS improves vMotion performance and multiple streams of IPv4/IPv6 TCP/UDP/IPSEC bandwidth over single interface between the Virtual Machines. For further information, see the User Manual section “ <i>NetQ RSS</i> ”.
<b>Default Queue RSS (DRSS)</b>	Allows the user to configure multiple hardware queues backing up the default RX queue. DRSS improves performance for large scale multicast traffic between hypervisors and Virtual Machines interfaces. For further information, see the User Manual section “ <i>Default Queue Receive Side Scaling (DRSS)</i> ”.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 13</a>
Rev. 4.15.5.10/4.5.5.10	
<b>General</b>	Added support for up to 8 ConnectX-4 Lx ports and up to 16 VFs. For further information, refer to the User Manual
<b>SR-IOV</b>	Added SR-IOV support for ESXi 5.5
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 13</a>
Rev. 4.15.4.1100/4.5.2.1100	
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 13</a>
Rev. 4.15.4.1000/4.5.2.1000	
<b>RX/TX Ring Resize</b>	Allows the network administrator to set new RX/TX ring buffer size.
Rev. 4.15.3.1008/4.5.1.1008	
<b>VXLAN Hardware Stateless Offloads for ConnectX®-4</b>	VXLAN hardware offload enables the traditional offloads to be performed on the encapsulated traffic.

**Table 8 - Change Log History (Sheet 2 of 2)**

Feature/Change	Description
<b>SR-IOV (Beta level)</b>	Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus. <b>Note:</b> SR-IOV is currently supported in ESXi 6.0 only.
<b>NetDump</b>	Enables a host to transmit diagnostic information via the network to a remote netdump service, which stores it on disk. Network-based core-dump collection can be configured in addition to or instead of disk-based coredump collection.
<b>NetQueue</b>	NetQueue is a performance technology in VMware ESXi that significantly improves performance in 10 Gigabit Ethernet virtualized environments.
<b>Wake-on-LAN</b>	Allows a network administrator to remotely power on a system or to wake it up from sleep mode